arXiv:1402.6306v1 [cond-mat.stat-mech] 25 Feb 2014

Superposition Enhanced Nested Sampling

Stefano Martiniani[†],[*] Jacob D. Stevenson,[†] David J. Wales, and Daan Frenkel
University Chemical Laboratories, University of Cambridge, Lensfield Road, Cambridge CB2 1EW, UK
(Dated: February 26, 2014)

The theoretical analysis of many problems in physics, astronomy and applied mathematics requires an efficient numerical exploration of multimodal parameter spaces that exhibit broken ergodicity. Monte Carlo methods are widely used to deal with these classes of problems, but such simulations suffer from a ubiquitous sampling problem: the probability of sampling a particular state is proportional to its entropic weight. Devising an algorithm capable of sampling efficiently the full phase space is a long-standing problem. Here we report a new hybrid method for the exploration of multimodal parameter spaces exhibiting broken ergodicity. Superposition enhanced nested sampling (SENS) combines the strengths of global optimization with the unbiased/athermal sampling of nested sampling, greatly enhancing its efficiency with no additional parameters. We report extensive tests of this new approach for atomic clusters that are known to have energy landscapes for which conventional sampling schemes suffer from broken ergodicity. We also introduce a novel parallelization algorithm for nested sampling.

Computer simulations play an important role in the study of phase transitions and critical phenomena. In particular, stochastic techniques such as Monte Carlo (MC) methods have proved to be powerful tools[1]. These methods rely on the ability of the Monte Carlo algorithm to sample the accessible volume in phase space. There are, however, many situations where standard Monte Carlo simulations suffer from a lack of ergodicity. In that case, more sophisticated algorithms are needed to explore the volume in phase space that is, in principle, accessible. Some such techniques are based on the efficient exploration of the underlying, multi-dimensional potential energy surface (PES)[2]. The PES, or energy landscape, can be viewed as a collection of basins separated by barriers, where each basin corresponds to a particular local minimum-energy configuration. The basin volumes define the entropic weight of the corresponding local minima. The transition rate from one basin to another depends on the barrier height as well as the relative entropic weights (configurational space volumes)[2]. Many PES of interest exhibit frustration in the form of low-lying minima with different morphologies separated by high barriers. These structures may act as kinetic traps, when fixed-temperature sampling methods such as molecular dynamics or Metropolis Monte Carlo sampling are used. There exist a wide range of extended or biased sampling techniques, both in Monte Carlo and in Molecular Dynamics, that make it possible to speed up the sampling of landscapes with kinetic traps. These techniques include Monte Carlo methods such as umbrella sampling[3, 4], density of states based methods, such as the Wang-Landau method [5] and replica exchange methods[6, 7], and their Molecular Dynamics counterparts. Examples are the replica-exchange MD method [8], and the Meta-Dynamics method [9]. In cases where a biased distribution is generated, the original distribution can be reconstructed using reweighting techniques[10, 11] However, these approaches may perform poorly when dealing with high-dimensional spaces exhibiting broken ergodicity or, in other words, with highly multimodal (or multifunnel[2, 12–14]) parameter spaces[15–18].

In recent years a Bayesian method known as nested sampling[19] has emerged as a possible alternative to extended or biased sampling methods. The nested-sampling approach has found widespread application in astrophysics[20, 21] and cosmology[22, 23], and has drawn the attention of computational and statistical physicists[24–29]. Furthermore the method has been recently adopted for Bayesian model comparison in system biology [30–32]. Nested sampling explores phase space in an unbiased way, and allows one to determine statistically the density of states associated with shrinking fractions of phase space. This objective is achieved by placing a constraint on the potential energy (for instance), which decreases at each nested sampling iteration. Like Wang-Landau sampling, the method is athermal and produces the density of states and the partition function (Bayesian evidence) as its primary product. However, nested sampling does not require binning of the energy for systems with continuous potentials. The self-adapting steps in energy (but constant in log phase space volume) is attractive because the approach does not require prior knowledge of possible phase transitions. For example, the step size adjusts automatically as the phase space volume shrinks near a first order phase transition[19, 25].

An important drawback of nested sampling is that when the decreasing energy constraint forbids a transition to an unexplored basin, that basin cannot be visited and ergodicity is broken. Hence, while nested sampling certainly is conceptually interesting, its performance is often no better than that of conventional extended sampling methods in dealing with systems exhibiting broken ergodicity[25]. In the present work we introduce a novel hybrid methodology for the exploration and thermodynamic analysis of such systems.

Superposition enhanced nested sampling (SENS) com-

bines the strengths of unbiased global optimization techniques [2] with those of nested sampling. Global optimization techniques such as basin-hopping [33, 34] are designed to find the lowest energy configuration of a PES. They are not constrained to sample according to any distribution, so they are free to use 'quick and dirty' techniques while searching for the global minimum. For example, they can take Monte Carlo steps that do not satisfy detailed balance, or make use of minimisation algorithms such as L-BFGS and conjugate gradient. Such operational freedom tends to make global optimization algorithms more efficient than generalised ensemble methods at locating the lowest energy minima[35–39]. Collections of the lowest energy minima configurations thus obtained can then be used in the context of the superposition approach (SA) [2, 38–44] to compute the thermodynamic properties of the system. However, doing so accurately at high temperatures using the SA alone often requires a prohibitively large number of minima.

In the present contribution we show how knowledge of the lowest energy minima and their statistical weights, calculated using the harmonic superposition approximation (HSA), can be exploited to enhance the problematic low energy behaviour of nested sampling, thus making it likely that none of the important minima and associated regions are missed. Although we discuss SENS in the context of energy landscapes, the method is completely general and can be applied to any multi-modal parameter space whose minima (maxima in likelihood) can be identified by global optimization algorithms.

NESTED SAMPLING

Nested sampling[19] provides an elegant solution to the problem of evaluating the density of states, and hence the partition function, for arbitrary systems. A likelihood value is assigned to each possible configuration. For our purposes the likelihood is the Boltzmann factor $\exp(-E/k_BT)$, but it could be some other measure. Typically, there are large numbers of configurations with a low likelihood. In addition, there may be a small number of configurations with high likelihood.

The aim of nested sampling is to sample configuration space uniformly, but with the energy constrained to lie below a maximum value, E^{max}, that decreases iteratively throughout the calculation. The rate of decrease is maintained self-consistently, such that the phase space volume with energy less than E^{max} decreases by a constant factor in each iteration.

The nested sampling algorithm starts by generating K configurations of the system completely at random, distributed uniformly, in configuration space. The energy, $E_{\mathcal{R}}$, of each of these configurations is computed and added to a sorted list. For each of these replicas we define the configurational phase space volume, $\Omega_{E\leq E_{\mathcal{R}}}$, containing all configurations with $E \leq E_{\mathcal{R}}$. The key insight of nested sampling is that the volumes, $\Omega_{E\leq E_{\mathcal{R}}}$, normalised by the total phase space volume, are distributed according to the Beta distribution, $\mathrm{Beta}(K-\mathcal{R}+1,\mathcal{R})$. This distribution has expectation value and variance

$$\mu_{\mathcal{R}} = 1 - \frac{\mathcal{R}}{K+1} \quad \text{and} \quad \sigma^2_{\mathcal{R}} = \frac{\mathcal{R}(K-\mathcal{R}+1)}{(K+2)(K+1)^2}. \tag{1}$$

The above formalism assumes that the total phase space volume Ω_{tot} is finite, but this condition can generally be satisfied with negligible error, for example by placing the system in a large box.

At the i-th nested sampling iteration the replica (out of K replicas) with highest energy E_i^{max} is discarded and replaced by a new configuration sampled uniformly under the constraint $E \leq E_i^{max}$. The maximum energy E_i^{max} is stored for later analysis. Again, the volume of configuration space with energy less than the $\mathcal{R}$-th largest energy, $\Omega_{E\leq E_{\mathcal{R}}}$, this time normalized by $\Omega_{E\leq E_i^{max}}$, is distributed according to the Beta distribution with mean and variance given by Eq. (1). During the nested sampling iteration the volume of phase space with energy below E^{max} contracts, on average, by $\mu_1 = K/(K+1)$. The algorithm produces a list of the form $\{E_1^{max}, E_2^{max}, \ldots, E_N^{max}\}$. We can associate a fraction of configuration space $X_i = \Omega_{E\leq E_i^{max}}/\Omega_{tot} = \mu_1^i$, with each E_i^{max}. The density of states, or the (normalized) volume of phase space with energy between E_{i+1}^{max} and E_i^{max} is simply

$$g_i(E) = X_i - X_{i+1} = \mu_1^i - \mu_1^{i+1} = \frac{1}{K+1}\left(\frac{K}{K+1}\right)^i. \tag{2}$$

Thermodynamic quantities of interest, such as the mean energy, entropy, free energy, and heat capacity, can easily be computed from the density of states at arbitrary temperature.

To generate configurations uniformly in space we use the strategy suggested by Skilling[19]: after removing the configuration with highest energy one of the remaining $K-1$ replicas (chosen randomly) is duplicated. The new configuration is then evolved through a Markov chain Monte Carlo (MCMC) walk sufficiently long to decorrelate the system from its initial state. This Monte Carlo walk is equivalent to a normal Monte Carlo simulation at infinite temperature. The coordinates are randomly perturbed, and the new configuration is accepted subject only to the condition that the energy remains below E^{max}. For most systems of interest the vast majority of the computational effort will be spent generating new configurations.

Parallelization

Nested sampling can be formulated to run in parallel on an arbitrary number of processors. We present a pseudocode description of our parallel implementation in Algorithm 2. Since this scheme also constitutes the basic framework for SENS we define the MCMC loop in the most general way at line 9 of Algorithm 2. For the purpose of discussing the algorithm in its simplest form here we will consider Algorithm 1.

for $l = 0$ to N-steps **do**
    generate trial configuration (e.g. by random perturbation);
    if $E_{trial} \leq E_{max}$: accept trial configuration;
end for

ALGORITHM 1. Nested sampling `MCLoop`

At each nested sampling iteration, instead of removing only the replica with the highest energy, we remove the $\mathcal{P}$ replicas with highest energy, where $\mathcal{P}$ is the number of processors available. The rate of phase space contraction now is given by $\mu_{\mathcal{P}}$, leading to much faster phase space contraction and shorter calculations in terms of wall-clock time. This parallelization procedure was first described in reference[27]. Our improvement is that we do not discard the $\mathcal{P} - 1$ replicas with highest energy but we store them for later analysis. Phase space contraction between iterations is still constant, but now the post-analysis is slightly more complicated. The fraction of configuration space associated with the n-th recorded energy is

$$X_n = \prod_{i=0}^{n} \frac{K - i\%\mathcal{P}}{K + 1 - i\%\mathcal{P}}, \tag{3}$$

where "%" is the mod operator. This method follows the same stepping routine as the existing parallelization algorithm. However it produces $\mathcal{P}$ times as many points, hence providing a more detailed picture of the potential energy surface and a much more fine-grained binning of the density of states.

SENS - THE CONCEPT

Global optimization is a common numerical problem and global optimization algorithms have been developed in many areas of science[2, 45, 46]. Knowledge of the local minima alone, however, is not sufficient to infer all the observables properties of interest from the energy landscape (or in general any parameter space). The harmonic superposition approximation (HSA) [47] (for more details, see e.g. [2]) allows one to compute density of states and partition function, solely based on the knowledge of the individual local minima and the local curvatures (normal mode frequencies) of the potential energy landscape, via the Hessian matrix. In the HSA each local minimum corresponds to a harmonic basin and observable properties are expressed as a sum over individual contributions of the minima.

▷ initialisation
1: generate K random configurations;
2: store their coordinates and energy in LiveList ;
▷ main loop
3: **while** termination condition **is False do**
4: remove the $\mathcal{P}$ replicas $\{\mathbf{R}_m^{(1)}, \ldots, \mathbf{R}_m^{(\mathcal{P})}\} \equiv \{\mathbf{R}_m\}$ with highest energy $\{E_m^{(1)} > \cdots > E_m^{(\mathcal{P})}\} \equiv \{E_m\}$ from LiveList ;
5: append $\{E_m\}$ to OutputList ;
6: set $E_{max} = E_m^{(\mathcal{P})}$;
7: select $\mathcal{P}$ replicas $\{\mathbf{R}_s^{(1)}, \ldots, \mathbf{R}_s^{(\mathcal{P})}\} \equiv \{\mathbf{R}_s\}$ from LiveList at random;
8: add a copy of $\{\mathbf{R}_s\}$ to LiveList ;
9: `MCLoop`$\{\{\mathbf{R}_s\}, E_{max},$ minima.db $\}$
10: **end while**

ALGORITHM 2. Parallel nested sampling

The HSA has been shown to be very effective for several systems[18, 39, 48] but the accuracy depends on how well the potential energy of the basins can be approximated as harmonic, and how many minima are thermodynamically important. The HSA becomes an increasingly good approximation at lower energies. The total number of minima increases exponentially with system size[47, 49], but it is impossible to tell a-priori how many of those are important. For example, LJ_{31}, a cluster of 31 isotropic particles interacting through a Lennard-Jones potential[50], has about 3×10^{15} predicted distinct configurations [12], but only a few dozen are required to reproduce the low temperature thermodynamic behaviour.

The global resolution of nested sampling depends on the number of replicas, K, used in the simulation, which is generally limited by the available computation time (the larger K, the slower the descent in energy). A more serious problem for nested sampling is that if the barrier to enter an unexplored funnel or superbasin is higher than the energy constraint E_{max}, that region of the PES will never be explored if it is not already populated. For example, in a crystallisation transition, at high energy the statistical weight of the liquid phase will be overwhelming and there will be no replicas in the region corresponding to the solid phase. However, as the energy constraint decreases (hence the temperature) the relative statistical weight associated with the solid phase increases. If we could sample phase space uniformly then at low energy we would observe a phase transition corresponding to crystallisation, but we must resort to a MCMC walk to explore phase space. If the entrance to the superbasin corresponding to the crystal has been locked out by E_{max} a Markov chain will not be able to find it, thus missing the transition.

Here we propose a new method that combines complementary techniques: nested sampling can sample efficiently the high energy regions of phase space, while at low energy a database of minima obtained by global optimization is used to augment the survey. While nested sampling assigns the correct statistical weight to each basin, global optimization makes it likely that no important minima are missed. This philosophy is also used in other methods combining replica exchange Monte Carlo with global optimization algorithms to treat broken ergodicity [12, 16].

SENS - THE ALGORITHM

Employing knowledge of low-lying minima fits naturally within the framework of nested sampling. We present here an exact and an approximate implementation of the SENS algorithm. Exact SENS is fully unbiased and requires no additional parameters than those needed in nested sampling. Approximate SENS, on the other hand, is formally biased and requires additional parameters. The reason for presenting both methods is that, in some cases, the latter can be considerably simpler to implement than the former while generally producing equally good, or better, results. SENS is based on the original nested sampling algorithm presented in Algorithm 2. The novelty of our method resides in the augmented sampling of the parameter space obtained by coupling the MCMC to the HSA. SENS can therefore be implemented by changing the function `MCLoop`($\{\mathbf{R}_s\}$, E_{max}, `minima.db`) of Algorithm 2. A full outline of the SENS algorithm can be found in Algorithm 3 of the Supplementary Information. To run SENS, a database of the lowest energy minima must be precomputed. In this work we adopt basin-hopping[33, 34] as the global optimization algorithm of choice.

Exact SENS

An unbiased version of SENS can be implemented by means of Hamiltonian replica exchange Monte Carlo moves[51]: in addition to normal MC steps, we introduce a new Monte Carlo step in which a minimum is sampled from the database according to its HSA configurational entropic weight:

$$w_c^{(b)}(E) = \frac{\Omega_c^{(b)}(E)}{\Omega_c(E)}. \tag{4}$$

We define the configurational volume of basin b

$$\Omega_c^{(b)}(E) \propto \frac{n_b (E - V^{(b)})^{\frac{\kappa}{2}}}{\prod_{\alpha=1}^{\kappa} \nu_\alpha^{(b)}}, \tag{5}$$

and the total configurational volume

$$\Omega_c(E) \propto \sum_b \frac{n_b (E - V^{(b)})^{\frac{\kappa}{2}}}{\prod_{\alpha=1}^{\kappa} \nu_\alpha^{(b)}}, \tag{6}$$

where $V^{(b)}$ is the potential energy of the minimum corresponding to basin b, $\nu_\alpha^{(b)}$ are the normal mode vibrational frequencies defined by the Hessian matrix, κ is the number of vibrational degrees of freedom, and n_b is the degeneracy of the basin (for Lennard-Jones clusters this is the number of distinct non-superimposable permutation-inversion isomers for minimum b)[2]. Here we have left out all the constant factors that cancel out and overall rotations. Once a minimum is selected, a configuration with $E \leq E_{max}$ is then chosen uniformly from within its basin of attraction. This approach corresponds to selecting a point uniformly from a multidimensional harmonic well. Such a configuration can be generated analytically, see the Supplementary Information for details. Unlike Ref. [16], in our approach we sample from the uniform distribution of configurations below energy E_{max}, rather than from the corresponding canonical distribution.

Thus, we obtain a configuration $\mathbf{R}_{sys}$ sampled according to the true Hamiltonian, $\mathbb{H}_{sys}$, and a trial configuration, $\mathbf{R}_{har}$, sampled according to the HSA-Hamiltonian, $\mathbb{H}_{har}$. The energies of the two configurations are then computed with the other Hamiltonian. If

$$\mathbb{H}_{har}(\mathbf{R}_{sys}) \leq E_{max} \text{ and } \mathbb{H}_{sys}(\mathbf{R}_{har}) \leq E_{max}, \tag{7}$$

then the true distribution and the HSA distributions overlap, the swap is accepted, and $\mathbf{R}_{sys}$ becomes $\mathbf{R}_{har}$, otherwise it is rejected. This procedure guarantees that detailed balance is satisfied. In practice only the lowest energy minima will successfully swap, since the HSA can only be reasonably accurate around these basins. It is, however, at low energy that such swaps are needed the most due to the hard energy constraint used by nested sampling. Note that swaps are complemented by regular MCMC walks to allow for the exploration of the full configuration space. In SENS the replicas are allowed to "tunnel" between basins, thus improving the sampling. A more detailed description, along with a pseudo-code implementation of `MCLoop` specific to Lennard-Jones clusters is provided in the Supplementary Information.

Approximate SENS

The implementation of approximate SENS is somewhat simpler, but comes at the cost of at least one extra parameter. The basic idea of approximate SENS is that the sampling of configuration space can be augmented by starting a MCMC walk from a local minimum configuration, sampled from the database according to its entropic weight Eq. (4), with some user defined frequency.

This frequency is intrinsically defined in exact SENS by the relative overlap of the HSA and the true density of states. To implement approximate SENS we only need a database of minima and their relative weights computed according to Eq. (4). Before each MCMC step a random number is drawn. If this number is less than some user defined probability, P_{DS}, then a minimum is selected from the database according to the HSA weights and the MCMC walk starts from this minimum configuration. A pseudo-code implementation of `MCLoop` for approximate SENS is provided in the Supplementary Information.

There are two main sources of bias in the approximate SENS. The first one is due to the limited number of minima from which we sample, since we cannot include the large number of high energy minima. The second source of error is due to the poor quality of the HSA approximation far from the minimum, hence the entropic weights for the minima are not accurate at high energy. The most obvious way of reducing these biases is to use long MCMC walks. In fact, if we sample from the wrong basin a long MCMC walk will allow the system to escape and explore regions of phase space with greater entropic weight. However, very long MCMC walks are computationally expensive, and if short runs are required we need to sample from the database of minima carefully. If we start sampling from the database of minima at high energy we will possibly introduce a bias due to over-weighting of the low energy regions of configuration space. To avoid this problem we suppress sampling from the database until we are sure the HSA is likely to describe the potential energy landscape accurately. We use a simple function (of the Fermi type) that delays the onset of sampling from the database of minima and limits its maximum frequency

$$f_{onset} = \frac{f_{max}}{1 + e^{(E_{min}^{(\mathcal{R})} - E_{on})/\Delta E}}, \tag{8}$$

where E_{on} is some onset energy and $E_{min}^{(\mathcal{R})}$ is the energy of the replica with lowest energy. E_{on} could be chosen as $E_{max}^{(\mathsf{minima.db})}$, the energy of the highest known minimum (stored in the database), or as the largest energy at which the HSA describes the system accurately. f_{max} and ΔE are user-defined parameters that determine the total probability of sampling after the onset and the width of the onset region, respectively. For small sampling probabilities, $P_{DS} \ll 1$, the optimal frequency of sampling from the database, should scale as $1/K$; a theoretical justification is derived in the Supplementary Information. Hence, for $P_{DS} \ll 1$, we can make the probability of sampling from the database independent of the number of replicas, replacing f_{max} with f_{max}/K.

We identify two possible strategies for the application of approximate SENS. One is to start sampling from a large database early in the simulation when $E_{on} = E_{max}^{(\mathsf{minima.db})}$, with a small P_{DS}, hence we choose $f_{max} \ll 1$. This procedure allows nested sampling to do most of the work, but ensures that no important basins will be missed. Alternatively, sampling from the database can be delayed until all the high temperature transitions have occurred, at which point we start sampling more extensively from the database, hence $f_{max} \gtrsim 1/2$. Note that the database can be considerably smaller in this case. The first strategy is a slight enhancement to nested sampling, while the latter strategy interpolates between nested sampling and the HSA in a similar spirit to the basin sampling method[12]. Importantly, even if we sample from the database of minima, we use the MCMC walk to explore more the anharmonic regions of a basin, allowing us to go beyond the harmonic approximation.

RESULTS

We test SENS by calculating the thermodynamic properties of Lennard-Jones clusters exhibiting broken ergodicity. Lennard-Jones (LJ) clusters are systems of particles that interact via the Lennard-Jones potential[50]

$$E = 4\epsilon \sum_{i<j} \left[\left(\frac{\sigma}{r_{ij}}\right)^{12} - \left(\frac{\sigma}{r_{ij}}\right)^{6} \right], \tag{9}$$

where ϵ is the pair well depth, σ is the separation at which $E = 0$, and $2^{1/6}\sigma$ is the equilibrium pair separation. LJ clusters have served as benchmarks for many global optimization techniques and for the estimate of thermodynamic properties by computer simulation[2, 12, 15, 18, 25, 34].

The majority of putative ground states for LJ clusters are based on icosahedral packings[13]. For some *magic number* LJ clusters complete Mackay icosahedra are possible, for examples $N = 13, 55$. Complete icosahedral structures are considerably more stable than neighbouring sizes and their landscape is funneled towards the global minimum[13]. There are, however, other sizes for which the global minimum is not icosahedral. Examples are LJ_{38}, whose ground state is an fcc-truncated octahedron[13], and LJ_{75} whose global minimum is a Marks decahedron[13]. Due to the overwhelming number (entropic weight) of structures based on incomplete icosahedra at high energy, the energy landscapes of LJ clusters with nonicosahedral global minima exhibit broken ergodicity. Calculating accurate thermodynamic properties for these systems has proved to be a real challenge for all conventional techniques[2, 15, 18, 25] and hybrid or more complicated schemes[12, 14, 15, 18] are necessary. LJ clusters with broken ergodicity therefore provide excellent benchmarks to test the performance of new sampling techniques.

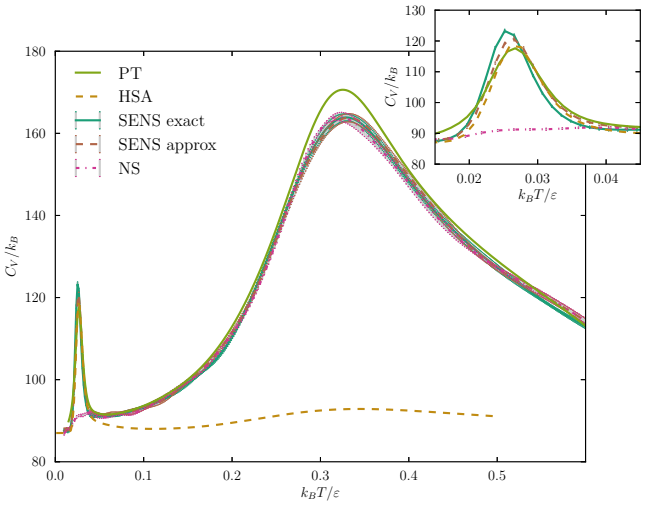

FIG. 1. Heat capacity curves for LJ_{31}. PT and HSA correspond to parallel tempering and the harmonic superposition approximation, respectively. All SENS calculations were performed using $K = 20000$ replicas.

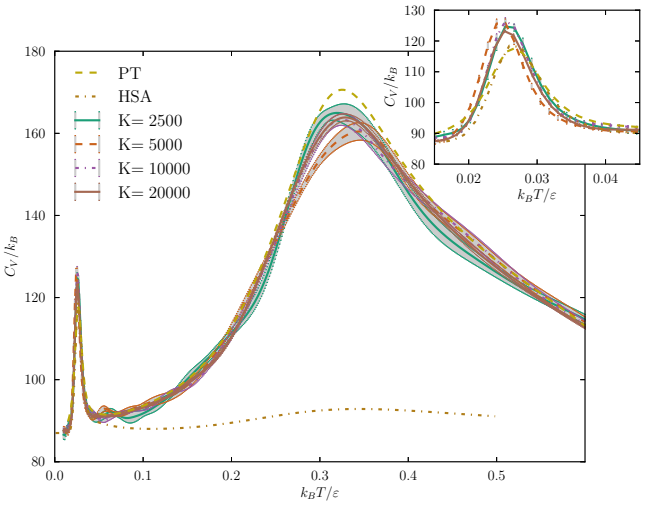

FIG. 2. Comparison of heat capacity curves for LJ_{31} obtained by exact SENS using different numbers of replicas. The PT and HSA curves were obtained by parallel tempering and the harmonic superposition approximation, respectively.

LJ_{31}

LJ_{31} is the smallest Lennard-Jones cluster exhibiting broken ergodicity and a low temperature solid-solid phase-like transition from Mackay to anti-Mackay surface structures[13]. Convergence of the heat capacity curve for LJ_{31} by parallel tempering (PT) with 24 geometrically distributed temperatures in the range 0.0125 to 0.6 required $N_E^{total} = 1.9 \times 10^{11}$ energy evaluations to converge (curve shown in Fig. 1). Partay et al.[25] report that $K = 288000$ replicas and $N_E^{\rm total} = 3.4 \times 10^{12}$ energy evaluations were needed to converge the heat capacity curve of LJ_{31} by nested sampling (NS) using a low particle density of $2.31 \times 10^{-3}\sigma^{-3}$ (100 fold less dense than our system). Fig. 1 compares the heat capacity curves obtained by PT, HSA (computed using $\gtrsim 80000$ minima), NS and SENS for LJ_{31}. The SENS and NS results correspond to $K = 20000$ replicas, $N = 10000$ steps for each MCMC walk, and $\mathcal{P} = 16$ cores. The database of minima used for SENS contained the lowest 183 minima although for SENS exact we observe that only seven minima contribute to the swaps; see Table IV of Supplementary Information for the swap statistics. From Fig. 1 we see that both exact SENS and approximate SENS are well converged and agree with the PT curve over the whole temperature range, and with the HSA at low temperature. We note that $K = 20000$ replicas are not nearly enough for NS to converge, and the low temperature peak is in fact completely absent. Using this number of replicas SENS requires half the total number of energy evaluations of PT and one order of magnitude less than NS, see Table I. The swap operations do not constitute a noticeable overhead and the reduction in the total number of energy evaluations corresponds to an equivalent reduction in wall-clock time.

LJ_{31}			
Method	K	N	$N_{\rm E}^{(total)}$
PT			1.9×10^{11}
NS ref.[25]	280000		3.4×10^{12}
NS	20000	10000	1×10^{11}
SENS approx	20000	10000	1×10^{11}
SENS exact	20000	10000	1×10^{11}
SENS exact	10000	10000	5.2×10^{10}
SENS exact	5000	10000	2.6×10^{10}
SENS exact	2500	10000	1.3×10^{10}

TABLE I. Comparison of methods used to obtain the LJ_{31} heat capacity curves shown in Figs. 1 and 2. $N_{\rm E}^{(total)}$ indicates the total number of energy evaluations (summed over all processors). PT was performed using 24 replicas spread geometrically through the temperature range 0.0125 to 0.6. Note that approximate SENS can perform as well as exact SENS when fewer replicas are used, in the interest of brevity we do not include these results as the LJ_{75} calculations illustrate clearly the capabilities of the method.

In Fig. 2 we show a comparison of PT, HSA and exact SENS for a range of replica numbers $2500 \leq K \leq 20000$; see Table I for comparison. We observe that the high temperature peak practically converges for $K = 10000$ and it resembles the features of the converged curve quite well even for smaller numbers of replicas. The low temperature peak instead converges very quickly, for as few as $K = 2500$ replicas, representing an improvement in performance of 20 times over PT. We note that one of the great strengths of SENS is that even when a small number of replicas are used and run times are very short,

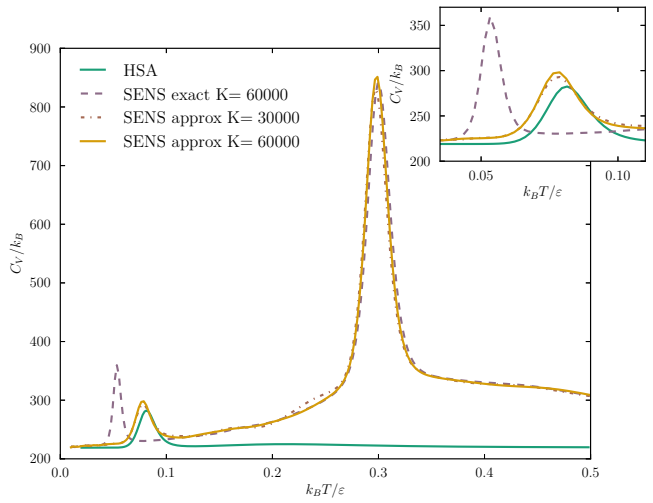

FIG. 3. Heat capacity curves for LJ_{75}. The PT and HSA results were obtained by parallel tempering and the harmonic superposition approximation, respectively. Exact SENS calculations were performed using $K = 60000$ replicas, while results for approximate SENS calculations are shown for both $K = 30000$ and $K = 60000$ replicas.

although the curves may not be completely converged, the physical picture produced by the method is always correct because all the important basins are visited. On the other hand, rapid convergence of the heat capacity curves, requires the HSA to be a good representation for the system. LJ_{38} is an example for which this condition does not hold, see the Supplementary Information for further details.

LJ_{75}

LJ_{75} is a particularly clear example of a double-funneled energy landscape[13] with $O(10^{25})$ distinct local minima[12]. The decahedral global minimum is separated by a very large potential energy barrier from the lowest icosahedral minimum. Sharapov and Mandelshtam[16] showed that $O(10^{12})$ (total) energy evaluations of adaptive parallel tempering are not enough to converge the heat capacity peak corresponding to the solid-solid phase-like transition in LJ_{75}[16]. Furthermore, the rate of convergence slows down dramatically (it practically stops) after $O(10^{11})$ (total) energy evaluations and coupling of PT to the HSA is necessary to obtain convergence of the low temperature peak[16]. Fig. 3 compares the heat capacity curves obtained by HSA (computed using 758 minima) and SENS for LJ_{75}. SENS was carried out using $K = 30000$ or $K = 60000$ replicas and $N = 10000$ steps for each MCMC walk on $\mathcal{P} = 16$ processors. The database of minima for SENS contained the lowest 758 minima. Approximate SENS started sampling from the database at $E_{on} = -369\ \epsilon$, while for exact

LJ_{75}			
Method	K	N	$N_{\mathrm{E}}^{(total)}$
SENS approx	30000	10000	4×10^{11}
SENS approx	60000	10000	8×10^{11}
SENS exact	60000	10000	8×10^{11}

TABLE II. Comparison of methods used to obtain the LJ_{75} heat capacity curves shown in Fig. 3. $N_{\mathrm{E}}^{(total)}$ indicates the total number of energy evaluations (summed over all processors). PT curves are not shown as the computational cost to converge its heat capacity by this method is computationally prohibitive as shown in Ref. [16]. SENS exact does not converge as quickly as approximate SENS due to the low accuracy of the HSA and hence the low swap acceptance.

SENS only 10 of the minima contributed to the swaps; see Table VI of the Supplementary Information for swap statistics. Unlike adaptive PT[16], approximate SENS converges in $O(10^{11})$ energy evaluations (Table II), but exact SENS fails to converge the low temperature peak for the same number of replicas. As for LJ_{38}, exact SENS does not converge quickly due to the lower accuracy of the HSA, as inferred from the extremely low swap acceptance (Table VI of Supplementary Information). On the other hand approximate SENS performs considerably better than for LJ_{38} because the melting transition is well separated from the solid-solid transition, thus allowing sampling from the database relatively early on in the simulation (right after melting) without affecting the melting transition.

Methods

We define a move in MCMC as the displacement of each individual particle along a random vector (n in total). After each MCMC walk we update the step size in order to keep the average acceptance ratio within range of some target value, which we have chosen as 0.5. The default parameter values for the onset function Eq. (8) were $f_{max} \approx 2/3$ and $\Delta E = 1$. We used a spherical box of radius $R = 2.5\ \sigma$ for $n = 31$, $R = 2.8\ \sigma$ for $n = 38$ and $R = 3.0\ \sigma$ for $n = 75$, with no periodic boundary conditions and no cutoff radius. All calculations were carried out on a single workstation with $\mathcal{P} = 16$ processors (eight-core dual Xeon E5-2670 2.6GHz, Westmere) using the improved parallelization scheme discussed in sec. Parallelization. The calculations were terminated when the energy difference between the replicas with highest and lowest energies was less than $10^{-2}\ \epsilon$. Energies of the final "live" replicas were added to the output and the compression factor associated with the ℓ^{th} "live" replica was computed as

$$\mu_\ell^{(live)} = \prod_{j=0}^{\ell<K} \frac{K-j}{K-j+1}. \qquad (10)$$

Error bars were obtained by the compression factor resampling scheme discussed in the Supplementary Information. By nested sampling or SENS iterations, $N_{\rm iter}$, we mean a whole nested sampling iteration on $\mathcal{P}$ processors, the total number of energy evaluations is $N_E^{(tot)} = N \times \mathcal{P} \times N_{\rm iter}$, where N is the number of steps in a MCMC. The computational overhead associated with global optimization by Basin Hopping is insignificant as $\leq O(10^5)$ energy evaluations are necessary to find the global minima of the LJ clusters considered here[52]. Highly modular Python/C parallel implementations of nested sampling and SENS are publicly available[53, 54].

We gratefully acknowledge discussions with Andy Ballard, Victor Ruehle, Thomas Stecher, Tine Curk, Boris Fackovec and Robert Baldock. This work has been supported by the European Research Council and the EPSRC. S.M. acknowledges financial support from the Gates Cambridge Scholarship and J.D.S. from the Marie Curie IEF 275544. D.F. acknowledges support from the ERC Advanced Grant 227758, Wolfson Merit Award 2007/R3 of the Royal Society of London.

* sm958@cam.ac.uk; Corresponding author
† These authors have contributed equally to this work

[1] D. Landau and K. Binder, *A guide to Monte Carlo simulations in statistical physics*, 3rd ed. (2009)
[2] D. J. Wales, *Energy Landscapes: Applications to Clusters, Biomolecules and Glasses* (2003) p. 681
[3] G. M. Torrie and J. P. Valleau, Journal of Computational Physics **23**, 187 (1977)
[4] B. Berg and T. Neuhaus, Physical Review Letters **68**, 9 (1992)
[5] F. Wang and D. Landau, Physical Review Letters **86**, 2050 (2001)
[6] R. H. Swendsen and J.-S. Wang, Physical Review Letters **57**, 2607 (1986)
[7] R. Swendsen and J.-S. Wang, Physical Review Letters **58**, 86 (1987)
[8] Y. Sugita and Y. Okamoto, Chemical Physics Letters **314**, 141 (1999)
[9] A. Laio and M. Parrinello, Proceedings of the National Academy of Sciences **99**, 12562 (2002)
[10] A. Ferrenberg and R. Swendsen, Physical Review Letters **61**, 2635 (1988)
[11] J. D. Chodera, W. C. Swope, J. W. Pitera, C. Seok, and K. A. Dill, Journal of Chemical Theory and Computation **3**, 26 (2007)
[12] D. J. Wales, Chemical Physics Letters **584**, 1 (2013)
[13] J. P. K. Doye, M. A. Miller, and D. J. Wales, The Journal of Chemical Physics **111**, 8417 (1999)
[14] J. P. K. Doye and D. J. Wales, Physical Review Letters **80**, 1357 (1998)
[15] P. Poulain, F. Calvo, R. Antoine, M. Broyer, and P. Dugourd, Physical Review E **73**, 056704 (2006)
[16] V. A. Mandelshtam, P. A. Frantsuzov, and F. Calvo, The Journal of Physical Chemistry A **110**, 5326
[17] V. A. Sharapov and V. A. Mandelshtam, The journal of physical chemistry. A **111**, 10284 (2007)
[18] V. Sharapov, D. Meluzzi, and V. Mandelshtam, Physical Review Letters **98**, 105701 (2007)
[19] J. Skilling, Bayesian Analysis **1**, 833 (2006)
[20] P. Mukherjee, D. Parkinson, and A. R. Liddle, The Astrophysical Journal **638**, L51 (2006)
[21] F. Feroz and M. P. Hobson, Monthly Notices of the Royal Astronomical Society **384**, 449 (2008)
[22] J. R. Shaw, M. Bridges, and M. P. Hobson, Monthly Notices of the Royal Astronomical Society **378**, 1365 (2007)
[23] F. Feroz, M. P. Hobson, and M. Bridges, Monthly Notices of the Royal Astronomical Society **398**, 1601 (2009)
[24] I. Murray, D. MacKay, Z. Ghahramani, and J. Skilling, in *Advances in Neural Information Processing Systems 18*, edited by Y. Weiss, B. Schölkopf, and J. Platt (2005) pp. 947–954
[25] L. B. Pártay, A. P. Bartók, and G. Csányi, The Journal of Physical Chemistry B **114**, 10502 (2010)
[26] S. O. Nielsen, The Journal of Chemical Physics **139**, 124104 (2013)
[27] N. S. Burkoff, C. Várnai, S. A. Wells, and D. L. Wild, Biophysical journal **102**, 878 (2012)
[28] L. B. Pártay, A. P. Bartók, and G. Csányi, Physical Review E **89**, 022302 (2014)
[29] B. J. Brewer, L. B. Pártay, and G. Csányi, Statistics and Computing **21**, 649 (2010)
[30] S. Aitken and O. E. Akman, BMC Systems Biology **7**, 72 (2013)
[31] R. Dybowski, T. J. McKinley, P. Mastroeni, and O. Restif, PloS one **8**, e82317 (2013)
[32] N. Pullen and R. J. Morris, PloS one **9**, e88419 (2014)
[33] Z. Li and H. A. Scheraga, Proceedings of the National Academy of Sciences of the United States of America **84**, 6611 (1987)
[34] D. J. Wales and J. P. K. Doye, The Journal of Physical Chemistry A **101**, 5111 (1997)
[35] D. J. McGinty, The Journal of Chemical Physics **55**, 580 (1971)
[36] J. J. Burton, The Journal of Chemical Physics **56**, 3133 (1972)
[37] M. R. Hoare, Adv. Chem. Phys. **40**, 49 (1979)
[38] F. H. Stillinger and T. A. Weber, Science **225**, 983 (1984)
[39] B. Strodel and D. J. Wales, Chemical Physics Letters **466**, 105 (2008)
[40] D. J. Wales, Molecular Physics **78**, 151 (1993)
[41] J. P. K. Doye and D. J. Wales, The Journal of Chemical Physics **102**, 9673 (1995)
[42] J. P. K. Doye and D. J. Wales, The Journal of Chemical Physics **102**, 9659 (1995)
[43] F. Calvo, J. P. K. Doye, and D. J. Wales, The Journal of Chemical Physics **115**, 9627 (2001)
[44] I. Georgescu and V. A. Mandelshtam, The Journal of Chemical Physics **137**, 144106 (2012)
[45] B. Hartke, Wiley Interdisciplinary Reviews: Computational Molecular Science **1**, 879 (2011)
[46] D. J. Wales, Science **285**, 1368 (1999)
[47] F. H. Stillinger and T. A. Weber, Science (New York, N.Y.) **225**, 983 (1984)
[48] S. Somani and D. J. Wales, The Journal of Chemical Physics **139**, 121909 (2013)
[49] D. J. Wales and J. P. K. Doye, The Journal of Chemical Physics **119**, 12409 (2003)

[50] J. E. Lennard-Jones, Proceedings of the Physical Society **43**, 461 (1931)
[51] H. Fukunishi, O. Watanabe, and S. Takada, The Journal of Chemical Physics **116**, 9058 (2002)
[52] M. T. Oakley, R. L. Johnston, and D. J. Wales, Physical Chemistry Chemical Physics **15**, 3965 (2013)
[53] S. Martiniani and J. D. Stevenson, "Nested Sampling https://github.com/js850/nested_sampling," (2013)
[54] S. Martiniani and J. D. Stevenson, "SENS https://github.com/smcantab/sens," (2013)

Supplementary Information

Sampling configurations in a harmonic well

Given a harmonic potential the configurational density of states for a basin can be obtained by inverse Laplace transforming the configurational partition function. In particular, the scaling goes as:

$$g_c(E) = \mathcal{L}^{-1}\left\{Z_c(\beta)\right\} \propto (E_c - V)^{\frac{\kappa}{2}-1}, \tag{11}$$

where all the terms that do not depend on energy have been left out. $g_c(E)$ is the configurational density of states, Z_c the configurational partition function (evidence), $\beta = 1/k_B T$ is the inverse temperature, E_c the configurational energy, V the potential energy of the corresponding minimum and κ is the number of degrees of freedom (for a n-atoms cluster $\kappa = 3n-6$). We can write the energy of the simple harmonic oscillator as

$$E_c = V + \frac{1}{2}\xi r^2 \tag{12}$$

where r is the magnitude of the displacement vector and ξ is the stiffness of the harmonic spring. We want to determine the probability distribution of the configurational energy as a function of the displacement vector norm, $\xi^{\frac{1}{2}} r$, to perform analytical uniform sampling in the harmonic well. The unnormalised probability of finding a configuration between E_c and $E_c + \mathrm{d}E_c$ must be proportional to the configurational density of states, from Eq. (11):

$$p(E_c)\mathrm{d}E_c \propto (E_c - V)^{\frac{\kappa}{2}-1}\mathrm{d}E_c. \tag{13}$$

Denoting $q = \xi^{\frac{1}{2}} r$, by a simple change of variables we can express the energy probability distribution in terms of q:

$$p(E_c)\mathrm{d}E_c = p(E_c(q))J\mathrm{d}q \propto (q^2)^{\frac{\kappa}{2}-1} q\mathrm{d}q, \tag{14}$$

where the Jacobian $J = \mathrm{d}E_c/\mathrm{d}q = q$ and the equality simplifies to the probability density function

$$p(q)\mathrm{d}q \propto q^{\kappa-1}\mathrm{d}q. \tag{15}$$

Hence q must be distributed according to the power law cumulative distribution function $P(q) = q^{\kappa}$ (denoted $\mathrm{Pow}(\kappa)$) to obtain the correct distribution of energies.
In order to sample uniformly below some energy constraint E_{max} we first generate a random $\kappa-$dimensional vector $\boldsymbol{v}$ with norm $v \sim \mathrm{Pow}(\kappa) \in (0, 1]$ in the unit hypersphere. Then starting from Eq. (12) we write

$$\boldsymbol{q}_{usc} = \sqrt{\frac{2(E_{max} - V)}{\xi}}\boldsymbol{v} \tag{16}$$

where $\boldsymbol{q}_{usc}$ is the uniformly sampled configuration vector with energy E_c. It can easily be verified that $\boldsymbol{q}_{usc}$ has the correct inner product:

$$E_c = \frac{1}{2}\xi q_{usc}^2 = (E_{max} - V)v^2, \tag{17}$$

where again $v \sim \mathrm{Pow}(\kappa) \in (0, 1]$.

Onset function scaling

How well would SENS perform in a multifunneled landscape? Let us assume that SENS is using a database that stores the lowest minima m_a and m_b of two different funnels, with associated energies V_a and V_b, respectively. Assuming that funnel F_b has already been missed (F_b is not populated and E_{max} is lower than the lowest transition state that leads to F_b) and all replicas are already in funnel F_a, the probability of successfully sampling in F_b is then

$$\Pr(\text{success}|\text{SENS with } K \text{ replicas in } F_a) = 1 - \prod_{i=1}^{n}(1 - P_{DS}p_b(i))^{\mathcal{P}}, \tag{18}$$

where n is the number of iterations before the calculation terminates, $\mathcal{P}$ is the number of processors, P_{DS} is a user defined probability of sampling from the database, and $p_b(i)$ is the discrete probability density that the minimum sampled from the database is m_b. Assuming that the funnels are harmonic we obtain Eq. (19)

where $\nu_\alpha = \omega_\alpha/(2\pi)$ is the vibrational frequency of mode α and for an object corresponding to a point group with o independent symmetry operations, there are o permutation-inversion operations associated with barrierless reorientations[2]. The ideal value for P_{DS} must then satisfy an identity of the form

$$1 - \prod_{i=1}^{n(V_a)}(1 - P_{DS}p_b(i))^{\mathcal{P}} = \phi, \tag{20}$$

where $n(V_a)$ is the average number of steps necessary for descending from E_g to V_a and ϕ is a probability close to 1, say $\phi = 0.999$. Taking the logarithm of both sides, Eq. (20) can be rewritten as

$$\sum_{i=1}^{n(V_a)} \log(1 - P_{DS}p_b(i)) = \frac{\log(1-\phi)}{\mathcal{P}}. \tag{21}$$

For small P_{DS} linearisation then leads to

$$P_{DS} = -\frac{\log(1-\phi)}{\mathcal{P}\sum_{i=1}^{n(V_a)} p_b(i)}, \tag{22}$$

which provides an optimal value for P_{DS}. To calculate the average number of steps necessary to descend a harmonic basin, first we calculate an expression for the set

$$p_b(i) = \frac{\Omega_c^{(b)}}{\Omega_c^{(a)} + \Omega_c^{(b)}} = \frac{(E_i - V_b)^{\frac{\kappa}{2}}\Theta(E_i - V_b)\prod_{\alpha=1}^{\kappa}\nu_\alpha^{(a)} o_a}{(E_i - V_b)^{\frac{\kappa}{2}}\Theta(E_i - V_b) o_a \prod_{\alpha=1}^{\kappa}\nu_\alpha^{(a)} + (E_i - V_a)^{\frac{\kappa}{2}}\Theta(E_i - V_a) o_b \prod_{\alpha=1}^{\kappa}\nu_\alpha^{(b)}}. \tag{19}$$

of energies that would be obtained by nested sampling if at each step the configurational space was compressed exactly by μ. We note that

$$\mu = \frac{\Omega(E_{i+1})}{\Omega(E_i)} = \frac{(E_{i+1} - V)^{\frac{\kappa}{2}}}{(E_i - V)^{\frac{\kappa}{2}}} \tag{23}$$

from which we find

$$E_i = (E_0 - V)\mu^{\frac{2i}{\kappa}} + V, \tag{24}$$

where $\mu = 1 - \mathcal{P}/(K+1)$. The number of steps $n(V_a)$ necessary to descend from E_g to $E - V_a = 10^{-\delta}$ can be obtained by rearranging Eq. (24) to give

$$n(V_a) = -\frac{\kappa}{2}\frac{\delta\log(10) + \log(E_g - V_a)}{\log(\mu)}. \tag{25}$$

Finally, substituting Eq. (24) for E_i in Eq. (19) we can evaluate Eq. (20) numerically to obtain an optimal value for P_{DS} (approximating $n(V_a)$ to the nearest integer).
In the main text we introduce the onset function Eq. (8) and suggest that for small P_{DS} an optimal way to make the probability of sampling from the database of minima independent of the number of replicas, is to use the prefactor $1/K$, hence the maximum frequency to sample from the minima, should be f_{max}/K, where f_{max} is a user-defined parameter. Eq. (25) gives the average number of steps necessary to descend a particular harmonic basin as a function of the number of replicas and by substituting this result in Eq. (20) we can evaluate P_{DS}, the optimal probability of sampling from the database of minima. In Fig. we plot $f_{max} = KP_{DS}$ vs K. For large K the optimal value of sampling from the minima scales as $1/K$, thus justifying the use of the prefactor in the onset function.

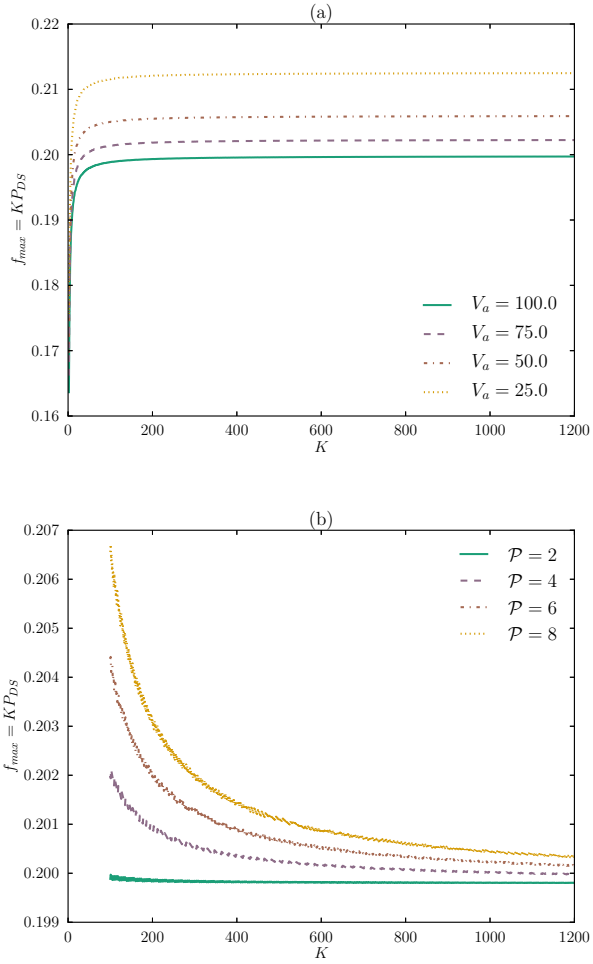

FIG. 4. Eq. (25) gives the average number of steps necessary to descend a particular harmonic basin as a function of the number of replicas K and by substituting it in Eq. (20) we can evaluate P_{DS}, the optimal probability of sampling from the database of minima. As long as K is sufficiently large, $f_{max} = KP_{DS}$ will be approximately constant, hence the optimal P_{DS} scales as $1/K$. (a) As V_a, the potential of m_a, decreases, the volume of F_a increases, and hence the optimal f_{max} increases as well. (b) The optimal f_{max} should be independent of the number of processors used, hence for sufficiently large K all curves approach the same value. Unless specified assume $\mathcal{P} = 1$, $\kappa = 3$, $E_g = 500$, $V_a = 100$, $V_b = 0$, $v_\alpha^{(a)} = 1$, $v_\alpha^{(b)} = 10$, $o_b = o_a = 1$, $\delta = 10$, $\phi = 0.999$.

Algorithms

SENS

A complete pseudo-code implementation of the SENS algorithm is provided in Algorithm 3.

▷ initialisation, set $i = 0$
generate a database of minima **in** minima.db ;
for minimum **in** minima.db **do**
    compute the Hessian matrix and its eigenvalues (needed to compute the HSA weight);
end for
while $i < K$ **do**
    sample a random configuration of the system;
    store its coordinates and its energy in LiveList ;
end while
▷ main loop
while termination condition **is False do**
    remove the $\mathcal{P}$ replicas $\{\mathbf{R}_m^{(1)}, \dots, \mathbf{R}_m^{(\mathcal{P})}\} \equiv \{\mathbf{R}_m\}$ with highest energy $\{E_m^{(1)} > \cdots > E_m^{(\mathcal{P})}\} \equiv \{E_m\}$ from LiveList ;
    append $\{E_m\}$ to OutputList ;
    set $E_{max} = E_m^{(\mathcal{P})}$;
    sample $\mathcal{P}$ replicas $\{\mathbf{R}_s^{(1)}, \dots, \mathbf{R}_s^{(\mathcal{P})}\} \equiv \{\mathbf{R}_s\}$ from LiveList at random;
    add a copy of $\{\mathbf{R}_s\}$ to LiveList ;
    MCLoop $\{\{\mathbf{R}_s\}$, minima.db , $E_{max}\}$;
end while
append LiveList to OutputList ;

ALGORITHM 3. Superposition Enhanced Nested Sampling

Exact SENS

A pseudo-code implementation of the MCLoop function for exact SENS can be found in Algorithm 4.

for all $\mathbf{R}_s$ **in** $\{\mathbf{R}_s\}$ **do**
    perform the Hamiltonian Exchange by evaluating $\mathbb{H}_{sys}(\mathbf{R}_{har}) = E_{sys}^{(\mathbf{R}_{har})}$ and $\mathbb{H}_{har}(\mathbf{R}_s) = E_{har}^{(\mathbf{R}_s)}$
    if $E_{har}^{(\mathbf{R}_s)} \leq E_{max}$ **and** $E_{sys}^{(\mathbf{R}_{har})} \leq E_{max}$ **then**
        swap $\mathbf{R}_{har} \leftrightarrow \mathbf{R}_s$;
        for $l = 0$ **to** N-steps **do**
            walk $\mathbf{R}_s$ by sampling uniformly within $\{E_s\} \leq E_{max}$;
        end for
    else reject the swap and perform a standard MCMC;
    end if
end for

ALGORITHM 4. Exact SENS MCLoop

The energy of a configuration $\mathbf{R}_s$ with respect to $\mathbf{R}_{min}$ is

$$E_{har}^{(\mathbf{R}_s)} = E_{har}^{(\mathbf{R}_{min})} + (\mathbf{R}_s - \mathbf{R}_{min})^\top \mathcal{H}_{min} (\mathbf{R}_s - \mathbf{R}_{min}), \tag{26}$$

where $\mathcal{H}_{min}$ is the Hessian matrix associated with $\mathbf{R}_{min}$. Note that more careful considerations are needed when dealing with Lennard-Jones clusters, in fact $\mathbf{R}_{min}$ is fixed in the database (hence also $\mathcal{H}_{min}$), thus breaking the translational and rotational invariance of the system. In order to avoid this problem we can either align $\mathbf{R}_{min}$ to $\mathbf{R}_s$, but then $\mathcal{H}_{min}$ must also be recalculated, or we can do the opposite, which is the most efficient solution. A pseudo-code implementation of the MCLoop function specific to this system is provided in Algorithm 5.

▷ this loop is performed in parallel
for all $\mathbf{R}_s$ **in** $\{\mathbf{R}_s\}$ **do**
    take the energy of $\mathbf{R}_s$ according to its LJ-Hamiltonian: $\mathbb{H}_{LJ}(\mathbf{R}_s) = E_{LJ}^{(\mathbf{R}_s)}$;
    sample a minimum m from minima.db according to its HSA entropic weight;
    analytically generate a configuration $\mathbf{R}_{har}$ in the harmonic well of m, with energy $\mathbb{H}_{har}(\mathbf{R}_{har}) = E_{har}^{(\mathbf{R}_{har})}$;
    evaluate $\mathbb{H}_{LJ}(\mathbf{R}_{har}) = E_{LJ}^{(\mathbf{R}_{har})}$;
    if $E_{LJ}^{(\mathbf{R}_{har})} > E_{max}$ **then**
        reject the exchange and perform a standard MCMC;
    else
        quench $\mathbf{R}_s \to \mathbf{R}_{que}$ with energy $\mathbb{H}_{LJ}(\mathbf{R}_{que}) = E_{LJ}^{(\mathbf{R}_{que})}$ and find the minimum configuration $\mathbf{R}_{min}$ with the corresponding energy in minima.db ;
        find the set of permutations ($\mathcal{P}$) and rotations ($\mathcal{R}$) such that $\mathcal{P}\mathcal{R}\mathbf{R}_{que} = \mathbf{R}_{min}$;
        generate $\mathbf{R}_s^{(\mathcal{P}\mathcal{R})} = \mathcal{P}\mathcal{R}\mathbf{R}_s$, which is now aligned to $\mathbf{R}_{min}$;
        compute $\mathbb{H}_{har}(\mathbf{R}_s^{(\mathcal{P}\mathcal{R})}) = E_{har}^{(\mathbf{R}_s^{(\mathcal{P}\mathcal{R})})}$;
        if $E_{har}^{(\mathbf{R}_s^{(\mathcal{P}\mathcal{R})})} > E_{max}$ **then**
            reject the move and perform a standard MCMC;
        else *swap* $\mathbf{R}_{har} \leftrightarrow \mathbf{R}_s$;
            for $l = 0$ **to** N-steps **do**
                walk $\mathbf{R}_s$ by sampling uniformly within $\{E_s\} \leq E_{max}$;
            end for
        end if
    end if
end for

ALGORITHM 5. Exact SENS MCLoop for LJ-clusters

Approximate SENS

A pseudo-code implementation of the MCLoop function for approximate SENS is presented in Algorithm 6.

```
for all R_s in {R_s} do
    if u ~ Uniform(0,1) < P_DS then
        swap R_b ↔ R_s;
        for l = 0 to N-steps do
            walk R_s by sampling uniformly within
            {E_s} ≤ E_max;
        end for
    else reject the swap and perform a standard MCMC;
    end if
end for
```

ALGORITHM 6. Approximate `MCLoop`

Statistical Uncertainty by Compression Factor Resampling

The nested sampling algorithm produces as its primary product a list of parameters (in our case energies) with an associated fraction of configurational space $X_i = \prod_{j=0}^{i} t_j$, where the t_j are the compression factors sampled on a unit interval with probability distribution Beta$(K - \mathcal{P} + 1, \mathcal{P})$ and expectation value $\mu = 1 - \mathcal{P}/(K+1)$. The exploration of configuration space is the challenging and time consuming part of the algorithm, while the overhead due to the assignment of compression factors is almost irrelevant. In general we use the expectation value μ of these compression factors in order to find the bins of density of states, g, that we need to calculate thermodynamic properties. Given a set of energies obtained by nested sampling, the correct size of the bins for the density of states is one unique realisation of the compression factors t that we do not know a priori (it is for this reason that we use the expectation value μ). There is some statistical uncertainty associated with the distribution of the bin size, which is ultimately due to the distribution of compression factors t, which we know. Since we are interested in the distribution of some observable $Q(\boldsymbol{E})$, say the heat capacity, we can use a representative set of parameters $\widetilde{\boldsymbol{E}} = \widetilde{E}_1, \widetilde{E}_2, \ldots, \widetilde{E}_n$ (energies) obtained by nested sampling (the time consuming part) and sample c sets of compression factors $\boldsymbol{t}_\ell = t_1^{(\ell)}, t_2^{(\ell)}, \ldots, t_n^{(\ell)}$ to associate with this representative set of parameters. This procedure is justified by the fact that we are interested in the probability distribution of $Q(\widetilde{\boldsymbol{E}})$ given the joint probability distribution $p(\boldsymbol{t}_\ell) = p(t_1^{(\ell)})p(t_2^{(\ell)}) \ldots p(t_n^{(\ell)})$. The mean and variance of $Q(\widetilde{\boldsymbol{E}})$ are therefore

$$\left\langle Q(\widetilde{\boldsymbol{E}}) \right\rangle = \sum_{\ell=1}^{c} Q(\widetilde{\boldsymbol{E}}) p(\boldsymbol{t}_\ell) = \frac{1}{c} \sum_{\ell=1}^{c} Q(\widetilde{\boldsymbol{E}}|\boldsymbol{t}_\ell); \qquad (27)$$

$$\begin{aligned} Var\left(Q(\widetilde{\boldsymbol{E}})\right) &= \sum_{\ell=1}^{c} Q(\widetilde{\boldsymbol{E}})^2 p(\boldsymbol{t}_\ell) - \left(\sum_{\ell=1}^{c} Q(\widetilde{\boldsymbol{E}})^2 p(\boldsymbol{t}_\ell) \right)^2 \\ &= \frac{1}{c} \sum_{\ell=1}^{c} Q(\widetilde{\boldsymbol{E}}|\boldsymbol{t}_\ell)^2 - \left(\frac{1}{c} \sum_{\ell=1}^{c} Q(\widetilde{\boldsymbol{E}}|\boldsymbol{t}_\ell) \right)^2 . \end{aligned} \qquad (28)$$

A protocol for quantifying uncertainty of this form was already suggested by Skilling[19]. However, this method suffers from some flaws: it does not include an estimate of systematic uncertainties and does not include in any way an estimate of the uncertainty due to the incomplete sampling of configurational space. The latter effect is due to the fact that resampling is performed over one representative set of energies ($\widetilde{\boldsymbol{E}}$) that could be missing a whole part of configurational space.

Results

LJ_{38}

LJ_{38} has a double funnel energy landscape[42] and its heat capacity exhibits three peaks: the first high temperature peak corresponds to a vapour-liquid transition, the second peak corresponds to melting, and the low temperature peak corresponds to a solid-solid transition from the cuboctahedral global minimum to the lowest icosahedral minimum[14]. The PT heat capacity curve for LJ_{38} generally overestimates the temperature at which the solid-solid phase-like transition occurs and, consequently, disappears under the melting peak, and should appear, instead, as a small shoulder slightly under the melting peak[15, 17]. For this reason long calculations are necessary for the heat capacity to converge. Partay et al.[28] employed $K = 244000$ replicas and $O(10^{12})$ energy evaluations to resolve the heat capacity of LJ_{38} by NS. Fig. 5 compares the heat capacity curves obtained by PT (using a box of radius $R = 2.8\ \sigma$), HSA (computed using $\gtrsim 89000$ minima) and SENS for LJ_{38}. SENS was carried out using $K = 50000$ replicas and $N = 10000$ steps for each MCMC walk. The database of minima used for SENS contained approximately 89000 minima. It appears that neither version of SENS can outperform NS or parallel tempering, which require the same number of energy evaluations to converge; see Table III for comparison. Exact SENS fails due to the inaccuracy of the HSA, and Table V shows the number of effective swaps from and to the basin. These numbers are considerably smaller than for LJ_{31} (see Table IV), even if the total number of iterations for LJ_{38} is much larger. Approximate SENS should generally work independently of the quality of the HSA. Here, however, the three transitions overlap significantly thus preventing sampling from the database early

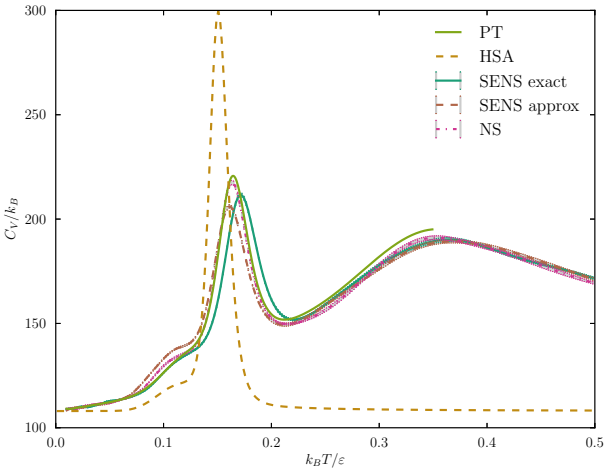

FIG. 5. Heat capacity curves for LJ_{38}. The PT and HSA results were obtained by parallel tempering and the harmonic superposition approximation, respectively.

	LJ_{38}		
Method	K	N	$N_{\mathrm{E}}^{(total)}$
PT			$O(10^{11})$
NS ref.[28]	244000		$O(10^{12})$
NS	50000	10000	3.3×10^{11}
SENS approx	50000	10000	3.3×10^{11}
SENS exact	50000	10000	3.3×10^{11}

TABLE III. Comparison of methods used to obtain the LJ_{31} heat capacity curves shown in Fig. 5. $N_{\mathrm{E}}^{(total)}$ indicates the total number of energy evaluations (summed over all processors).

enough to get the right low temperature behaviour without affecting the high temperature behaviour. This is not the case for LJ_{75} where the phase-like transitions are well separated in temperature and sampling can start early enough to get the correct low temperature behaviour without affecting the melting transition.

Swap statistics in exact SENS calculations

In this section we presents statistics for the exact SENS swaps in the longest runs of each LJ system presented in the paper. For each minimum we report the total number of swaps that led to this minimum, the number of effective swaps to this minimum (thus excluding the swaps within the same minimum) and the number of effective swaps that led from the minimum to another.

LJ_{31}			
minimum	total swaps to	effective swaps from	effective swaps to
−133.586421919	346376	4225	5994
−133.293821966	14585	7167	5403
−133.183574005	1529	1325	1325
−133.104620445	412	385	365
−132.998423589	25	11	25
−132.801757275	1	0	1
−132.765536037	0	1	0
−132.721370719	1	0	1

TABLE IV. Number of swaps to and from basins in exact SENS for LJ_{31} using $K = 20000$ replicas, $N = 10000$ steps for each MCMC and $\mathcal{P} = 16$. Total number of iterations per processor= 650025.

LJ_{38}			
minimum	total swaps to	effective swaps from	effective swaps to
−173.928426591	486584	245	48
−173.252378416	65	26	63
−173.134317009	163	27	154
−172.958633408	4	2	4
−172.877736411	38	6	37
−172.234926493	0	1	0
−171.992596189	1	0	1

TABLE V. Number of swaps to and from basins in exact SENS for LJ_{38} using $K = 50000$ replicas, $N = 10000$ steps for each MCMC and $\mathcal{P} = 16$. Total number of iterations per processor= 2082691.

LJ_{75}			
minimum	total swaps to	effective swaps from	effective swaps to
−397.492330983	632019	1	704
−396.282248826	15	296	12
−396.238512215	14	270	10
−396.193034959	7	41	5
−396.192994186	3	11	3
−396.191648856	4	45	4
−396.186860193	0	1	0
−396.126268882	2	63	1
−396.061061075	1	0	1
−396.061598578	0	13	0
−396.061061075	0	1	0
−396.057293139	1	0	1
−395.992783183	1	0	1

TABLE VI. Number of swaps to and from basins in exact SENS for LJ_{75} using $K = 60000$ replicas, $N = 10000$ steps for each MCMC and $\mathcal{P} = 16$. Total number of iterations per processor= 5043100.